IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

GUEST TEK INTERACTIVE ENTERTAINMENT LTD.,	)	
Plaintiff,	)	C.A. No. _____
v.	)	**JURY TRIAL DEMANDED**
NOMADIX, INC.,	)	
Defendant.	)	

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Guest Tek Interactive Entertainment Ltd. ("Guest Tek" or "Plaintiff"), for its Complaint against Defendant Nomadix, Inc. ("Nomadix" or "Defendant"), complains as follows:

NATURE OF THE ACTION

1. This is an action for infringement of United States Patent Nos. 8,811,184 ("the 184 Patent"), 9,154,435 ("the 435 Patent"), 9,531,640 ("the 640 Patent"), and 9,871,738 ("the 738 Patent (collectively, "the Patents-in-Suit").

THE PARTIES

2. Guest Tek is a corporation organized and existing under the laws of the province of Alberta, Canada with a place of business at Suite 600, 777 8 Ave., SW, Calgary, Alberta, T2P 3R5, Canada.

3. Upon information and belief, Nomadix is a corporation organized and existing under the laws of the State of Delaware, with its principal place of business at 30851 Agoura Road, Suite 102, Agoura, CA 91301.

JURISDICTION AND VENUE

4. This action arises under the patent laws of the United States, Title 35 of the United States Code. This Court has subject matter jurisdiction over this action under 28 U.S.C. §§ 1331 and 1338(a).

5. This Court has personal jurisdiction over Nomadix at least because it is incorporated under the laws of the State of Delaware.

6. Venue is proper in this judicial district under 28 U.S.C. §§ 1391(b), (c), and (d) and 1400(b) at least because Nomadix is incorporated under the laws of the State of Delaware.

BACKGROUND

7. On August 19, 2014, the United States Patent and Trademark Office ("USPTO") duly and legally issued the 184 Patent. Titled "Automatically Adjusting Bandwidth Allocated Between Different Zones In Proportion To The Number Of Users In Each Of The Zones Where A First-Level Zone Includes Second-Level Zones Not Entitled To Any Guaranteed Bandwidth Rate," the 184 Patent names David Ong as the inventor. A true and correct copy of the 184 Patent is attached to this complaint as Exhibit A.

8. On October 6, 2015, the USPTO duly and legally issued the 435 Patent. Titled "Automatically Adjusting Bandwidth Allocated Between Different Zones In Proportion To Summation Of Individual Bandwidth Caps Of Users In Each Of The Zones Where A First-Level Zone Includes Second-Level Zones Not Entitled To Any Guaranteed Bandwidth Rate," the 435 Patent names David T. Ong as the inventor. A true and correct copy of the 435 Patent is attached to this complaint as Exhibit B.

9. On December 27, 2016, the USPTO duly and legally issued the 640 Patent. Titled "Sharing Bandwidth Between Plurality Of Guaranteed Bandwidth Zones And A Remaining

Non-Guaranteed Bandwidth Zone," the 640 Patent names David T. Ong as the inventor. A true and correct copy of the 640 Patent is attached to this complaint as Exhibit C.

10. On January 16, 2018, the USPTO duly and legally issued the 738 Patent. Titled "Allocating Bandwidth Between Bandwidth Zones According To User Load," the 738 Patent names David T. Ong as the inventor. A true and correct copy of the 738 Patent is attached to this complaint as Exhibit D.

11. Guest Tek is the owner, by assignment registered in the USPTO, of the entire right, title, and interest in the Patents-in-Suit.

12. Guest Tek is a global leader in broadband technology and interactive solutions for the hospitality industry. Among Guest Tek's products are its OneView Internet and OneView Media solutions that allow hotels to offer internet and media services to their guests. These solutions are sometimes hereafter referred to collectively as the "OneView Products."

13. The OneView Products are covered by at least one claim of at least one of the Patents-in-Suit. Guest Tek has complied with the marking requirements of 35 U.S.C. § 287.

14. Guest Tek has incurred significant costs in connection with the research and development leading to the OneView Products, and in connection with its marketing of them. Guest Tek's One View Products have achieved substantial commercial success and won praise from the industry.

15. Upon information and belief, Defendant made, used, sold, or offered to sell in the United States, or imported into the United States, infringing instrumentalities such as the Nomadix Access Gateways (the "AG's"), which have been or are controlled by software called the Nomadix Service Engine ("NSE Software").

16. Upon information and belief, Defendant currently makes, uses, sells or offers to sell in the United States, or imports into the United States, infringing instrumentalities such as the AG's and NSE Software.

17. Upon information and belief, since 2016, Defendant has released several versions of the NSE Software for use in the AG's, which versions have included NSE Version 8.7, NSE Version 8.8, NSE Version 8.9, NSE Version 8.10, and NSE Version 8.11.

18. Upon information and belief, the structure, function and operation of NSE Version 8.7, NSE Version 8.8, NSE Version 8.9, NSE Version 8.10, and NSE Version 8.11 are not different from one another in material part, insofar as the structure, function and operation covered by the Patents-In-Suit are concerned.

19. Attached hereto as Exhibit E is a document entitled, "Nomadix Access Gateway User Guide, Version 8.8" (the "User Guide"), (obtained from http://www.nomadix.com/Websites/nomadix/download/agdocs/latest/AG-Userguide-8-8.pdf) which, upon information and belief, describes the structure, function and operation of at least NSE Version 8.8, as well as, insofar as the Patents-In-Suit are concerned, NSE Version 8.7, NSE Version 8.9, NSE Version 8.10, and NSE Version 8.11.

20. Attached hereto as Exhibit F is a document entitled, "Nomadix Advances the State of Bandwidth Management for Visitor-Based Networks, Guest Internet Providers to Benefit from New Capabilities Found in Version 8.7 of Nomadix Service Engine (the "Nomadix Press Release") (obtained from http://www.nomadix.com/Websites/nomadix/download/press-releases/87software.pdf) which, upon information and belief, describes the structure, function and operation of several features, including a share unused bandwidth feature for class-based

queuing, released in NSE Version 8.7, as well as, insofar as the Patents-In-Suit are concerned, NSE Version 8.8, NSE Version 8.9, NSE Version 8.10, and NSE Version 8.11.

21. Upon information and belief, since at least 2016, Defendant has made, used, sold, or offered for sale in the United States or imported into the United States (and continues to make, use, sell, or offer for sale in the United States or import into the United States) AG's known as the Nomadix Access Gateway 2400 (the "AG 2400"), the Nomadix Access Gateway 2500 ("the AG 2500"), the Nomadix Access Gateway 5600 (the "AG 5600"), the Nomadix Access Gateway 5800 (the "AG 5800"), and the Nomadix Access Gateway 5900 (the "AG 5900"), among others, that have NSE Version 8.7, NSE Version 8.8, NSE Version 8.9, NSE Version 8.10, and/or NSE Version 8.11installed thereon. Furthermore, since at least 2016, Defendant has installed, or caused to be installed, NSE Version 8.7, NSE Version 8.8, NSE Version 8.9, NSE Version 8.10, and/or NSE Version 8.11on the AG 2400's, AG 2500's, AG 5600's, AG 5800's and/or AG 5900's, among others, in the United States.

22. Upon information and belief, Defendant has made, used, sold, or offered for sale in the United States, or imported into the United States, AG's other than the AG 2400's, AG 2500's, AG 5600's, AG 5800's, or AG 5900's (hereinafter "the Other AG's") that have NSE Version 8.7, NSE Version 8.8, NSE Version 8.9, NSE Version 8.10, and/or NSE Version 8.11 installed thereon, or Defendant has installed on, or caused to be installed on, the Other AG's, NSE Version 8.7, NSE Version 8.8, NSE Version 8.9, NSE Version 8.10, and/or NSE Version 8.11. The identity of the Other AG's will be determined through discovery.

23. The AG 2400's, AG 2500's, AG 5600's, AG 5800's, AG5900's, and the Other AG's which have run, or are currently running, NSE Version 8.7, NSE Version 8.8, NSE

Version 8.9, NSE Version 8.10, and NSE Version 8.11 in the United States are referred to hereinafter as "the Accused Products."

24. Upon information and belief, Defendant has released NSE Software other than NSE Version 8.7, NSE Version 8.8, NSE Version 8.9, NSE Version 8.10, and NSE Version 8.11 ("the Other NSE Versions"), the identity of which will be determined through discovery. The Other NSE Versions, when installed on an AG, may cause infringement of one or more of the Patents-In-Suit.

25. Upon information and belief, Defendant has made, used, sold or offered for sale in the United States, or imported into the United States (or currently makes, uses, sells, or offers for sale in the United States, or imports into the United States), the Other AG's that have the Other NSE Versions installed thereon, or has installed (or caused to be installed) the Other NSE Versions on the Other AG's in the United States. The identity of such Other AG's with such Other NSE Versions will be determined through discovery. The Other AG's with such Other NSE Versions may infringe one or more of the Patents-In-Suit.

26. Through its making, using, selling or offering for sale in the United States, or importing into the United States, the Accused Products ("Defendant's Activities"), Defendant has infringed and continues to infringe the Patents-in-Suit. Discovery will reveal whether Defendant's making, using, selling or offering for sale in the United States, or importing into the United States, the AG's running the Other NSE Versions also infringes the Patents-in-Suit.

COUNT I
INFRINGEMENT OF THE 184 PATENT

27. Guest Tek incorporates each of the preceding paragraphs as if fully set forth herein.

28. Establishments such as hotels typically allow their guests and users in conference-room meetings to use their computer network to access the internet. Staff may also use the computer network to access the internet for internal purposes. Entertainment systems (e.g., AV, games, etc.) in the hotel guest rooms may also be connected to the computer network. Usually, the bandwidth made available to establishments by their internet service providers is fixed or limited and must be shared over the number of users.

29. The 184 Patent (Exhibit A) is directed to a system and method of allocating the available bandwidth among users.

30. Defendant has infringed and is infringing at least claim 1 of the 184 Patent either literally or under the doctrine of equivalents. For ease of reference, claim 1 of the 184 Patent is reproduced below as an example of Defendant's infringement and each paragraph thereof has been annotated with an identifying letter:

> A bandwidth management system for allocating bandwidth between a plurality of bandwidth zones at an establishment serving a plurality of users, each zone having a number of users competing for bandwidth allocated to the zone, wherein the number of users in each zone changes over time; the bandwidth management system comprising:
>
> a) a plurality of queues, wherein each of the zones has a corresponding queue;
>
> b) an enqueuing module for receiving network traffic from one or more incoming network interfaces, determining a belonging zone to which the network traffic belongs, and enqueuing the network traffic on a queue corresponding to the belonging zone;
>
> c) a dequeuing module for selectively dequeuing data from the queues and passing the data to one or more outgoing network interfaces; and
>
> d) a quantum manager for dynamically adjusting values of a plurality of quantums, each of the queues having a respective quantum associated therewith;

e) wherein, when a selected queue has no guaranteed bandwidth rate or has already reached its guaranteed bandwidth rate, the dequeuing module dequeues at most an amount of data from the selected queue up to the quantum of the selected queue before dequeuing data from another of the queues;

f) the quantum manager dynamically adjusts the values of the quantums in proportion to the number of users in each of the zones as the number of users change over time;

g) the values of the quantums being automatically adjusted such that the quantum of a first queue is a higher value than the quantum of a second queue while the zone to which the first queue corresponds has a higher number of users than the zone to which the second queue corresponds, and such that the quantum of the first queue is a lower value than the quantum of the second queue while the zone to which the first queue corresponds has a lower number of users than the zone to which the second queue corresponds;

h) at least one of the zones is a first-level zone that includes a plurality of second-level zones not entitled to any guaranteed bandwidth rate;

i) network traffic enqueued on one or more queues corresponding to the second-level zones is dequeued and then enqueued on the queue corresponding to the first-level zone; and

j) the quantum manager determines the number of users of the first-level zone by accumulating the number of users under each of the second-level zones.

31. Page 11 of the User Guide describes that the Accused Products allocate bandwidth between a plurality of bandwidth zones (e.g., classes) at an establishment (e.g., a hotel property) serving a plurality of users. Pages 11 to 12 further describe how each zone (e.g., class) has a number of users competing for bandwidth allocated to the zone. Page 12 of the User Guide likewise describes how the number of users in each zone (e.g., class) changes over time. These pages reveal that the Accused Products comprise a bandwidth management system that allocates bandwidth between a plurality of bandwidth zones at an establishment serving a plurality of

users, each zone having a number of users competing for bandwidth allocated to the zone, wherein the number of users in each zone changes over time, and therefore satisfy the preamble of claim 1 of the 184 Patent.

32. Page 8 of the User Guide states that the NSE's core package includes Class-Based Queueing and pages 11 to 12 of the User Guide provide a "Use Case" and an "Example Illustration of Class-Based Queueing." As shown, there are three different classes (Conference, Guest Room, Public) that are utilized in a particular example of class-based queuing. On information and belief, each of the three classes has its own queue since these three classes are queued on a class basis. These portions of the User Guide reveal that the Accused Products provide a plurality of queues, wherein each of the zones has a corresponding queue, and therefore satisfy paragraph a) of claim 1 of the 184 Patent.

33. Page 13 of the User Guide illustrates how each class's network traffic is handled separately based on its corresponding queue's bandwidth level. Page 4 of the User Guide, for example, reveals that the Accused Products have network interfaces in the form of Ethernet ports. On information and belief, the Accused Products contain structure and/or functionality that provides an enqueuing module for receiving network traffic from one or more incoming network interfaces, determining a belonging zone to which the network traffic belongs (e.g., one of Lobby, VIP guests, or Meeting Room), and enqueuing the network traffic on a queue corresponding to the belonging zone. This structure and functionality satisfies paragraph b) of claim 1 of the 184 Patent.

34. The illustration at page 13 of the User Guide demonstrates that the Accused Products contain structure and functionality that removes data from each class's corresponding queue. On information and belief, this data is sent to an outgoing network interface such as the

WAN port. Therefore, the Accused Products contain structure and functionality that provides a dequeuing module for selectively dequeuing data from the queues and passing the data to one or more outgoing network interfaces. This structure and functionality satisfies paragraph c) of claim 1 of the 184 Patent.

35. The "Example Illustration of Class-Based Queueing" at page 12 of the User Guide demonstrates that the Accused Products change the bandwidth allocated to each class, and hence to each queue associated with each class, over time. On information and belief, the NSE assigns an amount of data that can be served at any time from a class and passed to another zone or to the internet, and that amount of data is determined or otherwise limited by a quantum. Because each class's effective bandwidth is changed, the Accused Products include structure or functionality to dynamically adjust the values of quantums for each queue thereby changing the bandwidth for each queue. The Accused Products thus provide a quantum manager for dynamically adjusting values of a plurality of quantums, each of the queues having a respective quantum associated therewith. This structure and functionality satisfies the paragraph d) of claim 1 of the 184 Patent.

36. On information and belief, the structure and functionality of an AG 2500 running NSE Version 8.9 is substantially the same as the structure and functionality of the other Accused Products, insofar as the Patents-In-Suit are concerned.

37. Page 103 of the User Guide demonstrates that minimum bandwidth settings are respected regardless of priority. The term "minimum bandwidth" represents the amount of bandwidth that will be made available when needed even if the WAN link is otherwise saturated. The throughput estimator running on a AG 2500 running NSE Version 8.9 ("Test Unit"), for example, shows that when there are two classes (e.g., Plan A and Other) that are competing for

bandwidth, where Plan A has already reached its minimum/guaranteed bandwidth rate of 500Kbps and Other has no minimum/guaranteed bandwidth rate, the dequeuing module within the tested AG 2500 provides effective throughput for Plan A as 600Kbps and for the Other class of 1430Kbps. On information and belief, in order to achieve this 600/1430 ratio of bandwidth amounts, the Accused Products cycle back and forth, dequeuing network traffic from each class, and the relative amounts of data dequeued from each class before switching to the other class together forms the same ratio in order to achieve the desired effective throughputs. On information and belief, the maximum amount of data dequeued from Plan A is limited to a quantum that is at most 630/1430 of the quantum of the Other queue in order to achieve the desired effective throughput of 630Kbps for Plan A and 1430Kpbs for Other. These descriptions in the User Guide, and the testing on the Test Unit, demonstrate that, in the Accused Products, when a selected queue has no guaranteed bandwidth rate or has already reached its guaranteed bandwidth rate, the dequeuing module dequeues at most an amount of data from the selected queue up to the quantum of the selected queue before dequeuing data from another of the queues. This structure and functionality satisfies paragraph e) of claim 1 of the 184 Patent.

38. The structure and functionality described under the heading "Example Illustration of Class-Based Queueing" on pages 12-13 of the User Guide demonstrates how the available bandwidth is allocated in proportion to the number of users in each of the zones as the users change over time. The quantum is a maximum amount of data dequeued from a queue before dequeuing data from another queue. Thus, the ratio of effective bandwidths of different zones corresponds to the ratio of their quantums. Since the effective throughput bandwidths outputted are controlled by the Accused Products to be in proportion to the number of users in each zone as the users change over time, on information and belief, the Accused Products include structure or

functionality to dynamically adjust the quantums in proportion to the number of users in each of the zones as the number of users change over time. On information and belief, the quantum manager controlling the quantums in this manner allows the Accused Products to control the resulting effective throughputs to be proportional to the number of users in each class. Furthermore, page 2 of the Nomadix Press Release discloses how the bandwidth allocations (and therefore the quantums) are adjusted in proportion to the subscribers when using the "share unused bandwidth" feature incorporated into the Accused Products. These pages of the User Guide and Nomadix Press Release demonstrate that the Accused Products provide a quantum manager that dynamically adjusts the values of the quantums in proportion to the number of users in each of the zones as the number of users change over time. This structure and functionality satisfies the paragraph f) of claim 1 of the 184 Patent.

39. A throughput simulator running on the Test Unit shows that an effective throughput of Plan A is higher than the effective throughput of the Other class when the user load on Plan A is higher than the load on the Other class. Likewise, testing of the Test Unit, configured with class-based queuing, weighted fair queuing, and share bandwidth all enabled, demonstrates that, while there are two users logged into a Plan A, a single user under Plan B loses some of her bandwidth. When Plan A has more users than Plan B, the aggregate bandwidth rate of Plan A is higher than Plan B. The situation reverses when Plan B has more users than Plan A. On information and belief, the Accused Products automatically adjust the quantums to thereby control the throughput speeds of both Plan A and Plan B as the numbers of users in the two classes change over time. On information and belief, the quantum (achieving 200Kbps) of a first queue (for Plan A) is a higher value than the quantum (achieving 133Kbps) of a second queue (for Plan B) while the zone (Plan A) to which the first queue corresponds has a higher

number of users (2 users v. 1 user) than the zone (Plan B) to which the second queue corresponds. Likewise, on information and belief, the relative sizes of the quantums reverse when Plan B has more users that Plan A. These tests and observable results reveal that the functionality of the Accused Products automatically adjusts the quantums such that the quantum of a first queue is a higher value than the quantum of a second queue while the zone to which the first queue corresponds has a higher number of users than the zone to which the second queue corresponds, and such that the quantum of the first queue is a lower value than the quantum of the second queue while the zone to which the first queue corresponds has a lower number of users than the zone to which the second queue corresponds. This structure and functionality satisfies paragraph g) of claim 1 of the 184 Patent.

40. Page 103 of the User Guide describes how first-level zones and second level zones can be configured. The terms "top-level classes" and "subclasses" as used in the User Guide equate to "first-level" and "second-level zones" as used in the patent claims. Page 104 of the User Guide shows a configuration where at least one of the classes is a first-level zone (PriorityOne) that includes a plurality of second-level zones (SubClassOne, SubClassTwo, SubClassThree). Page 104 of the User Guide shows how each class and subclass can have its attributes changed. As shown, one attribute that can be changed to any value including a zero (0) value is minimum upload / download speed. Minimum speed as described in the User Guide represents the guaranteed bandwidth rate for that class. Users are free to set the upload/download minimum guaranteed bandwidth rates to zero and to configure multiple zero bandwidth guarantee subclasses for each top-level class. Testing the class-based queuing configuration screen on the Test Unit shows a first-level zone (PriorityFour) that includes a plurality of second-level zones (FourSubOne, FourSubTwo, FourSubThree) not entitled to any guaranteed

bandwidth rate (i.e., FourSubOne, FourSubTwo, FourSubThree all have zero "Min" bandwidth specified). These descriptions and observations demonstrate that the Accused Products function such that at least one of the zones is a first-level zone that includes a plurality of second-level zones not entitled to any guaranteed bandwidth rate, and therefore satisfy paragraph h) of claim 1 of the 184 Patent.

41. Page 104 of the User Guide describes how the bandwidth of various sub-classes add up to form the total bandwidth assigned to a particular class. For instance, the effective throughput under a top-level class (e.g., Priority One, 30000Kbps) includes all the various network traffic from its subzones (e.g., SubClassOne, 15000Kbps + SubClassTwo, 10000Kbps + SubClassThree, 5000Kbps). In other words, the network traffic of each of these second-level classes is dequeued and then enqueued on the top level class for transmission. This description demonstrates that, in operation of the Accused Products, network traffic enqueued on one or more queues corresponding to the second-level zones is dequeued and then enqueued on the queue corresponding to the first-level zone. This structure and functionality satisfies paragraph i) of claim 1 of the 184 Patent.

42. Page 104 of the User Guide shows how user load under a particular top level class (e.g., Priority One, 84800Kbps) corresponds to all the user load numbers under the various second-level classes (SubClassOne, 35780Kbps + SubClassTwo, 29900Kbps + SubClassThree, 19120Kbps). This illustrates one example of how the Accused Products determine the load of the top-level zone by accumulating the user load under each of the second-level zones. Testing on the Test Unit shows changes to effective throughput for two zones (e.g., Plan A, Plan B) as the number of users in each zone changes over time. On information and belief, the Accused Products determine the number of users under a particular top-level class to be an accumulation

of all the users under each of the second-level zones. This is similar to how the offered load is accumulated in the throughput estimator example on page 104 of the User Guide. Page 103 of the User Guide states that subscribers can only be assigned to sub-classes. On information and belief, a subscriber added to a subclass will be accumulated under the corresponding top-level class. The dot notation utilized in the User Guide on page 105 of the User Guide makes it clear that the numbers of users and associated user load in all subclasses add up together and determine load of the top-level class. In particular, page 105 of the User Guide shows that subscribers can be added a priority1.subclass, which, by definition, means the user is a member of the priority1 class. This description demonstrates that in operation of the Accused Products, the quantum manager determines the number of users of the first-level zone by accumulating the number of users under each of the second-level zones. This structure and functionality satisfies paragraph j) of claim 1 of the 184 Patent.

43. Defendant has directly infringed, and continues to directly infringe, at least claim 1 of the 184 Patent by making, using, selling or offering for sale in the United States, or importing into the United States, the Accused Products, as described above.

44. Defendant has been on actual notice that it infringes the 184 Patent. Guest Tek has properly marked its products. Additionally, Defendant has had actual knowledge of its infringement at least since February 13, 2018 and February 22, 2018, when Guest Tek advised Defendant by letter of the infringement. Defendant has deliberately, intentionally and willfully infringed the 184 Patent, and continues to do so.

45. As a direct result of Defendant's infringing acts, Guest Tek has suffered and will continue to suffer damage and irreparable harm.

46. Unless Defendant and those acting in active concert with Defendant are enjoined from infringing the 184 Patent, Guest Tek will continue to suffer irreparable injury for which damages are an inadequate remedy.

COUNT II
INFRINGEMENT OF THE 435 PATENT

47. Guest Tek incorporates each of the preceding paragraphs as if fully set forth herein.

48. The 435 Patent (Exhibit B) is directed to allocating bandwidth over multiple zones.

49. Defendant has infringed and is infringing at least claim 19 of the 435 Patent either literally or under the doctrine of equivalents. For ease of reference, claim 19 of the 435 Patent is reproduced below as an example of Defendant's infringement and each paragraph thereof has been annotated with an identifying letter:

> A bandwidth management system for allocating bandwidth between a plurality of bandwidth zones at an establishment serving a plurality of users, each zone having a number of users competing for bandwidth allocated to the zone, wherein each of the users has an individual bandwidth cap, and at least one of the individual bandwidth caps changes over time, the system comprising:
>
> a) a plurality of queues, wherein each of the zones has a corresponding queue; and
>
> b) one or more processors configured to:
>
> c) receive network traffic from one or more incoming network interfaces;
>
> d) determine a belonging zone to which the network traffic belongs;
>
> e) enqueue the network traffic on a queue corresponding to the belonging zone;

f) selectively dequeue data from the queues; pass the data to one or more outgoing network interfaces; and

g) dynamically adjust values of a plurality of quantums, each of the queues having a respective quantum associated therewith;

h) wherein, when a selected queue has no guaranteed bandwidth rate or has already reached its guaranteed bandwidth rate, the one or more processors are configured to dequeue at most an amount of data from the selected queue up to the quantum of the selected queue before dequeuing data from another of the queues;

i) the one or more processors are configured to dynamically adjust the values of the quantums in proportion to a summation value of the individual bandwidth caps of the users in each zone as the individual bandwidth caps change over time;

j) the values of the quantums being automatically adjusted such that the quantum of a first queue is a higher value than the quantum of a second queue while the zone to which the first queue corresponds has a higher summation value of the individual bandwidth caps of the users than the zone to which the second queue corresponds, and

k) such that the quantum of the first queue is a lower value than the quantum of the second queue while the zone to which the first queue corresponds has a lower summation value of the individual bandwidth caps of the users than the zone to which the second queue corresponds;

l) at least one of the zones is a first-level zone that includes a plurality of second-level zones not entitled to any guaranteed bandwidth rate, and

m) the one or more processors are further configured to: dequeue network traffic enqueued on one or more queues corresponding to the second-level zones and then enqueue it on the queue corresponding to the first-level zone; and

n) determine the summation value of the individual bandwidth caps of the users of the first-level zone by accumulating the individual bandwidth caps of the users under each of the second-level zones.

50. The "Class-Based Queuing" section on Page 11 of the User Guide describes the Accused Products as a bandwidth management system for allocating bandwidth between a plurality of bandwidth zones (e.g., classes) at an establishment (e.g., property) serving a plurality of users. Pages 11 to 12 of the User Guide describe how each zone (e.g., class) has a number of users competing for bandwidth allocated to the zone. Page 11 of the User Guide describes how class-based queuing itself does not apply rules to individual users but states that bandwidth limits may still be applied to restrict individual users, if desired. The example of Weighted Fair Queuing on page 28 of the User Guide describes how users in different plans can have different individual bandwidth caps. Page 220 of the User Guide further describes how different plans can be set-up, where the plans each have individual bandwidth caps (in Up and Down directions) for each user and can place the users in desired classes, e.g., steps 7 and 8 in the Billing options set up process. Pages 222 to 232 of the User Guide describe various options of an Information and Control Console (ICC) provided by the Access Gateway (AG) that allows users to change their individual bandwidth caps. In particular, see page 222 of the User Guide stating that user can access the ICC popup screen to "select their bandwidth and billing plan options." Likewise, page 278 of the User Guide further explains that the ICC includes a "bandwidth selection" pull down menu that allows users to control and thereby change their individual bandwidth caps over time. These pages reveal that the Accused Products comprise a bandwidth management system for allocating bandwidth between a plurality of bandwidth zones at an establishment serving a plurality of users, each zone having a number of users competing for bandwidth allocated to the zone, wherein each of the users has an individual bandwidth cap, and at least one of the individual bandwidth caps changes over time, and therefore satisfy the preamble of claim 19 of the 435 Patent.

51. Page 8 of the User Guide shows the NSE core functionality is a bandwidth management system that includes Class Based Queueing. Pages 11 to 12 describes a plurality of classes. Since queuing is performed on a class-basis in class-based queuing and since there are a plurality of classes, on information and belief, there are a plurality of queues in the Accused Products. As shown on pages 11 to 12 of the User Guide, there are three different classes (Conference, Guest Room, Public) that are utilized in a particular example of class-based queuing. On information and belief, each of the classes has its own queue since these three classes are queued on a per class basis. These pages of the User Guide reveal that the Accused Products have a plurality of queues, wherein each of the zones has a corresponding queue, and therefore satisfy paragraph a) of claim 19 of the 435 Patent.

52. Page 341 of the User Guide states that the Nomadix Access Gateway is a computing device that runs software. On information and belief, the Accused Products are computing devices that run software and include one or more processors configured to perform various recited actions (addressed below) as specified by the software instructions. This page demonstrates that the Accused Products satisfy paragraph b) of claim 19 of the 435 Patent.

53. Page 13 of the User Guide describes how each class's received network traffic is processed in a group separate from traffic in other classes in a class-based queuing system in order to achieve different throughput outputs for each class. The product picture provided on page 3 of the User Guide illustrates several network interfaces such as WAN, Eth1, Eth2, etc. On information and belief, the processors of the Accused Products receive network traffic from one or more incoming network interfaces. These pages demonstrate that the Accused Products satisfy paragraph c) of claim 19 of the 435 Patent.

54. Page 13 of the User Guide describes how each class's network traffic is processed in a group separate from traffic in other classes in a class-based queuing system in order to achieve different throughput outputs for each class. On information and belief, the processors of the Accused Products determine a belonging zone to which the network traffic belongs (e.g., one of Lobby, VIP guests, or Meeting Room). This page demonstrates that the Accused Products satisfy paragraph d) of claim 19 of the 435 Patent.

55. Page 13 of the User Guide describes how each class's network traffic is processed in a group separate from traffic in other classes in a class-based queuing system in order to achieve different throughput outputs for each class. On information and belief, the processors of the Accused Products enqueue the network traffic on a queue corresponding to the belonging zone (e.g., on a queue for one of "Lobby", "VIP Guests", "Meeting Room" in the example on page 13). Pages 11 to 12 of the User Guide describe how classes are queued on a class-basis and therefore each class will have its own queue. These pages demonstrate that the Accused Products satisfy paragraph e) of claim 1 of the 435 Patent.

56. Page 13 of the User Guide describes how each class's network traffic is processed in a group separate from traffic in other classes in a class-based queuing system in order to achieve different throughput outputs for each class. On information and belief, the processors of the Accused Products selectively dequeue the network traffic from the queues and pass the data to one or more outgoing network interfaces. The chart on page 13 of the User Guide shows how network traffic throughput is formed by network traffic from the three different queues (e.g., for each of Lobby, VIP Guests, and Meeting Room). On information and belief, the data is removed data from each class's queue and passed to the output network interface(s) (e.g., WAN, Eth1,

etc.) for transmission to achieve the indicated throughputs. This page demonstrates that the Accused Products satisfy paragraph f) of claim 19 of the 435 Patent.

57. The "Example Illustration of Class-Based Queueing" at page 12 of the User Guide demonstrates that the Accused Products change the bandwidth allocated to each class, and hence to each queue associated with each class, over time. On information and belief, the NSE assigns an amount of data that can be served at any time from a class and passed to another zone or to the Internet, and that amount of data is determined or otherwise limited by a quantum. Because each class's effective bandwidth is changed, it follows that the processors of the Accused Products dynamically adjust the values of quantums for each queue thereby changing the bandwidth for each queue. The Accused Products thus dynamically adjust values of a plurality of quantums, each of the queues having a respective quantum associated therewith. This structure and functionality satisfies the paragraph g) of claim 19 of the 435 Patent.

58. Page 103 of the User Guide dislcoses that minimum bandwidth settings are respected regardless of priority. The term "minimum bandwidth" as utilized in the User Guide equates to the term "guaranteed bandwidth rate" utilized in the patent claims. It represents the amount of bandwidth that will be made available when needed even if the WAN link is otherwise saturated. The throughput estimator running on the Test Unit shows that when there are two classes (Plan A, Other) that are competing for bandwidth, where Plan A has already reached its minimum bandwidth of 500Kbps and Other has no minimum bandwidth, the dequeuing module within the Test Unit provides effective throughput for Plan A as 600Kbps and for the Other class of 1430Kbps. On information and belief, in order to achieve this 600/1430 ratio of bandwidth amounts, the Accused Products cycle back and forth dequeuing network traffic from each class, and the relative amounts of data dequeued from each class before

switching to the other class together form the same ratio in order to achieve the desired effective throughputs. On information and belief, the maximum amount of data dequeued from Plan A is limited to a quantum amount that is at most 630/1430 of the quantum amount of the Other queue in order to achieve the desired effective throughput of 630Kbps for Plan A and 1430Kpbs for Other. This page, and testing on the Test Unit, demonstrate that, in the Accused Products, when a selected queue has no guaranteed bandwidth rate or has already reached its guaranteed bandwidth rate, the dequeuing module dequeues at most an amount of data from the selected queue up to the quantum of the selected queue before dequeuing data from another of the queues, and therefore demonstrate that the Accused Products satisfy paragraph h) of claim 19 of the 435 Patent.

59. Pages 12 to 13 of the User Guide describe how the available bandwidth is allocated in proportion to user load under the zones as the users change over time. Page 28 of the User Guide describes how Weighted Fair Queuing allocates bandwidth in proportion to individual or group bandwidth limits. The example on page 28 shows that when there are 100 users each having 1M individual caps in a first group competing for bandwidth with 100 users each having 2M individual caps in a second group, the second group gets double the bandwidth than the first group. The ratio is 2:1 because the sum of the 100 * 2M individual caps in the premium plan is double the 100 * 1M in the basic plan. The quantum represents a maximum amount of data that will be dequeued from a queue before dequeuing data from another queue. Thus, the ratio of effective bandwidths of different zones corresponds to the ratio of their quantums. Since the effective throughput bandwidths outputted are controlled by the Accused Products to be in proportion to the summation value of the individual bandwidth caps of the users in each zone as the users change over time, on information and belief the quantums for

each queue are also dynamically adjusted by the Accused Products in proportion to the summation value of the individual bandwidth caps of the users in each zone as the number of users change over time. Furthermore, page 2 of the Nomadix Press Release announcing the release of the "Share unused" bandwidth feature that works in conjunction with class-based queuing mentions how the bandwidth allocations (and therefore the quantums) are proportionally increased. On information and belief, controlling the quantums in this manner is what allows the Accused Products to cause the resulting effective throughputs to be proportional to the sum of bandwidth caps under each group. These pages demonstrate that the Accused Products satisfy paragraph i) of claim 19 of the 435 Patent.

60. On information and belief, the Nomadix Press Release describes the operation of the Accused Product, in relevant part, insofar as the structure, function and operation of the Patents-In-Suit is concerned.

61. Pages 13 to 14 of the User Guide show an example of using Weighted Fair Queuing in conjunction with Class-based Queuing. The example shows how the aggregate throughput of the two classes is adjusted in proportion to the summation of individual bandwidth caps of users in each class. As shown, in the second part of the test when the class 1 users come online, the users under class 2 lose some of their bandwidth. When class 1 has higher summation value of the individual bandwidth caps of the users than class 2, the aggregate rate of class 1 is higher than class 2. The quantums controlling the bandwidth throughput speeds of both class 1 and class 2 queues change as the summations of individual bandwidth caps of the users in the two classes change over time. Specifically, the graph and example show the quantum of a first queue (for class 1) is a higher value than the quantum of a second queue (for class 2) while the zone to which the first queue corresponds (class 1) has a higher summation value of the

individual bandwidth caps of the users than the zone to which the second queue corresponds (class 2). These pages demonstrate that the Accused Products satisfy paragraph j) of claim 19 of the 435 Patent.

62. Pages 13 to 14 of the User Guide show an example of using Weighted Fair Queuing in conjunction with Class-based Queuing. The example shows how the aggregate throughput of the two classes is adjusted in proportion to the summation of individual bandwidth caps of users in each class. As shown, in the second part of the test when the class 1 users come online, the users under class 2 lose some of their bandwidth. When class 1 has higher summation value of the individual bandwidth caps of the users than class 2, the aggregate rate of class 1 is higher than class 2. The quantums controlling the bandwidth throughput speeds of both class 1 and class 2 change as the numbers of summations of individual bandwidth caps of the users in the two classes change over time. Specifically, this graph and example show the quantum of a first queue (for class 2) is a lower value than the quantum of a second queue (for class 2) while the zone to which the first queue corresponds (class 2) has a lower summation value of the individual bandwidth caps of the users than the zone to which the second queue corresponds (class 1). These depictions and functionality demonstrate that the Accused Products satisfy paragraph k) of claim 19 of the 435 Patent.

63. Pages 103 of the User Guide describes how first-level zones and second level zones can be configured. These terms “top-level classes” and “subclasses” as used in the User Guide equate to first-level and second-level zones. Page 104 of the User Guide shows a configuration where at least one of the classes is a first-level class (PriorityOne) that includes a plurality of second-level classes (SubClassOne, SubClassTwo, SubClassThree). Page 104 of the User Guide shows how each class and subclass can have its attributes changed. As shown, one

attribute that can be changed to any value including a zero (0) value is minimum upload / download speed. Minimum speed as described in the User Guide represents the guaranteed bandwidth rate for that class. Users are free to set the upload/download minimum guaranteed bandwidth rates to zero and to configure multiple zero bandwidth guarantee subclasses for each top-level class. Testing the class-based queuing configuration screen on the Test Unit shows a first-level zone (PriorityFour) that includes a plurality of second-level zones (FourSubOne, FourSubTwo, FourSubThree) not entitled to any guaranteed bandwidth rate (e.g., FourSubOne, FourSubTwo, FourSubThree all have zero "Min" bandwidth specified). These pages and testing demonstrate that the Accused Products function such that at least one of the zones is a first-level zone that includes a plurality of second-level zones not entitled to any guaranteed bandwidth rate, and therefore satisfy paragraph l) of claim 19 of the 435 Patent.

64. Page 104 of the User Guide describes how the bandwidth of various sub-classes add up to form the total bandwidth assigned to a particular class. For instance, the effective throughput under a top-level class (e.g., Priority One, 30000Kbps) includes all the various network traffic from its subzones (e.g., SubClassOne, 15000Kbps + SubClassTwo, 10000Kbps + SubClassThree, 5000Kbps). In other words, the network traffic of each of these second-level classes is dequeued and then enqueued on the top level class for transmission. This description demonstrates that, in operation of the Accused Products, network traffic enqueued on one or more queues corresponding to the second-level zones is dequeued and then enqueued on the queue corresponding to the first-level zone. This structure and functionality satisfies paragraph m) of claim 19 of the 435 Patent.

65. Page 104 of the User Guide shows how user load under a particular top-level class (e.g., Priority One, 84800Kbps) corresponds to all the user load numbers under the various

second-level classes (SubClassOne, 35780Kbps + SubClassTwo, 29900Kbps + SubClassThree, 19120Kbps). In other words, the Accused Products determine the load of the top-level zone by accumulating the user load under each of the second-level zones. Although the chart on page 104 does not show accumulating the individual bandwidth caps of the users under each of the second-level zones, as shown and described in the examples on pages 13 to 14 of the User Guide, the Accused Products change effective throughput based on the summations of the individual bandwidth caps of the users under each zone. On information and belief, the total number of individual bandwidth caps under a particular top-level class are an accumulation of all the individual bandwidth caps of the users under each of the second-level zones similar to how the offered load is accumulated in the above throughput estimator. Page 103 of the User Guide states that subscribers can only be assigned to sub-classes. On information and belief, a subscriber added to a subclass will be accumulated under the corresponding top-level class. The dot notation utilized in the User Guide on page 105 of the User Guide makes it clear that the amounts of users and associated user load in all subclasses add up together and determine load of the top-level class. In particular, page 105 of the User Guide shows that subscribers can be added a priority1.subclass, which be definition means the user is a member of the priority1 class. These pages demonstrate that, in operation of the Accused Products, the processors determine the summation value of the individual bandwidth caps of the users of the first-level zone by accumulating the individual bandwidth caps of the users under each of the second-level zones, and therefore the Accused Products satisfy paragraph n) of claim 19 of the 435 Patent.

66. Defendant has directly infringed, and continues to directly infringe, at least claim 19 of the 435 Patent by making, using, selling or offering for sale in the United States, or importing into the United States, the Accused Products, as described above.

67. Defendant has been on actual notice that it infringes the 435 Patent. Guest Tek has properly marked its products. Additionally, Defendant has had actual knowledge of its infringement at least since February 13, 2018 and February 22, 2018, when Guest Tek advised Defendant by letter of the infringement. Defendant has deliberately, intentionally and willfully infringed the 435 Patent, and continues to do so.

68. As a direct result of Defendant's infringing acts, Guest Tek has suffered and will continue to suffer damage and irreparable harm.

69. Unless Defendant and those acting in active concert with Defendant are enjoined from infringing the 435 Patent, Guest Tek will continue to suffer irreparable injury for which damages are an inadequate remedy.

COUNT III
INFRINGEMENT OF THE 640 PATENT

70. Guest Tek incorporates each of the preceding paragraphs as if fully set forth herein.

71. The system and method described and claimed in the 640 Patent (Exhibit C) allocates bandwidth over multiple zones.

72. Defendant has infringed and is infringing at least claim 1 of the 640 Patent either literally or under the doctrine of equivalents. For ease of reference, claim 1 of the 640 Patent is reproduced below and each paragraph thereof has been annotated with an identifying letter:

> A system comprising:
>
> a) one or more first network interfaces coupled to a first network with which a fixed total amount of bandwidth is available to transfer data;
>
> b) one or more second network interfaces coupled to a second network;

c) a plurality of queues, each of the plurality of the queues corresponding to a respective one of a plurality of bandwidth zones,

d) the plurality of the bandwidth zones including a plurality of first level guaranteed bandwidth zones and only one first level remaining bandwidth zone not entitled to any guaranteed bandwidth;

e) and one or more processors operable to:

f) determine a belonging bandwidth zone to which network traffic received from either of the first or the second network interfaces belongs;

g) enqueue the network traffic received from either of the first or the second network interfaces on a queue corresponding to the belonging bandwidth zone; and

h) cycle though the plurality of the queues, dequeue the network traffic from the plurality of the queues, and thereafter pass the network traffic dequeued from the plurality of the queues to one of the first or the second network interfaces for transmission to a destination network address;

i) wherein, when dequeing the network traffic from a particular queue, the one or more processors are operable to automatically determine an amount of the network traffic to dequeue from the particular queue according to a bandwidth limit for the particular queue;

j) the bandwidth limit for each of the plurality of the queues corresponding to the plurality of the first level guaranteed bandwidth zones includes a guaranteed bandwidth rate; and

k) the bandwidth limit for the first level remaining bandwidth zone has no guaranteed bandwidth rate but includes a bandwidth cap equal to the fixed total amount of bandwidth of the one or more first network interfaces minus the guaranteed bandwidth rate for each of the plurality of the first level guaranteed bandwidth zones.

73. Page 45 of the User Guide provides a guide for installing an AG for use by "system administrators" thereby indicating that the Accused Products provide a system. This page demonstrates that the Accused Products satisfy the preamble of claim 1 of the 640 Patent.

74. Page 38 of the User Guide shows that the Accused Products include one or more first network interfaces (WAN / Eth2) coupled to a first network (the Internet) with which a fixed total amount of bandwidth (100Mbps on WAN and 20Mpbs on Eth2) is available. Page 116 of the User Guide shows configuration of the fixed WAN port and Up/Down link speeds (Kbps) for network interfaces on the Access Gateway. Page 104 of the User Guide shows configuration of a fixed total amount of bandwidth (40000Kbps) available on the WAN port. Page 11 of the User Guide discloses another example where the WAN link has a fixed speed of 100Mbps link. Page 12 of the User Guide discusses a "saturated link of 200Mbps" indicating that the fixed total amount of 200Mbps in that example. Likewise, page 13 of the User Guide states another example where the WAN interface has a global upper limit of 900M. These pages of the User Guide reveal that the Accused Products include one or more first network interfaces coupled to a first network with which a fixed total amount of bandwidth is available to transfer data, and therefore the Accused Products satisfy paragraph a) of claim 1 of the 640 Patent.

75. Page 38 of the User Guide shows that the AG's includes one or more second network interfaces (Eth1) coupled to a second network (the Guest HSIA network). This page of the User Guide reveals that the Accused Products include one or more second network interfaces coupled to a second network, and therefore the Accused Products satisfy paragraph b) of claim 1 of the 640 Patent.

76. Page 8 of the User Guide states that the NSE's core package includes Class-Based Queueing and pages 11 to 12 of the User Guide provide a "Use Case" and an "Example Illustration of Class-Based Queueing." As shown, there are three different classes (Conference, Guest Room, Public) that are utilized in a particular example of class-based queuing. The Conference, Guest Room, Public classes are a plurality of bandwidth zones. On information and

belief, each of the three classes has its own queue since these three classes are queued on a class basis. These pages of the User Guide reveal that the Accused Products provide a plurality of queues, each of the plurality of the queues corresponding to a respective one of a plurality of bandwidth zones, therefore satisfying paragraph c) of claim 1 of the 640 Patent.

77. Page 104 of the User Guide shows the throughput simulator of a bandwidth zone configuration having a plurality of first level guaranteed bandwidth zones (PriorityOne, PriorityTwo, PriorityThree) and only one first level remaining bandwidth zone not entitled to any guaranteed bandwidth (Other). Page 14 of the User Guide further explains how, by default, only the eighth top-level class has no guaranteed bandwidth. Testing the class-based queuing configuration screen on the Test Unit shows a plurality of first level guaranteed bandwidth zones (PriorityOne, PriorityTwo, PriorityThree) and only one first level remaining bandwidth zone (PriorityFour) not entitled to any guaranteed bandwidth rate. These pages of the User Guide and the testing reveal that, in the Accused Products, the plurality of the bandwidth zones including a plurality of first level guaranteed bandwidth zones and only one first level remaining bandwidth zone not entitled to any guaranteed bandwidth, therefore satisfying paragraph d) of claim 1 of the 640 Patent.

78. Page 341 of the User Guide states that an AG is a computing device that runs software. On information and belief, the Accused Products are computing devices that include one or more processors operable to perform various actions as specified by software instructions, and therefore satisfy paragraph e) of claim 1 of the 640 Patent.

79. Page 13 of the User Guide describes how each class's network traffic is processed in a group separate from traffic in other classes in a class-based queuing system in order to achieve different throughput outputs for each class. On information and belief, the processors of

the Accused Products determine a belonging bandwidth zone to which network traffic received from either of the first or the second network interfaces belongs (e.g., one of Lobby, VIP guests, or Meeting Room). This page demonstrates that the Accused Products satisfy paragraph f) of claim 1 of the 640 Patent.

80. Page 13 of the User Guide describes how each class's network traffic is processed in a group separate from traffic in other classes in a class-based queuing system in order to achieve different throughput outputs for each class. On information and belief, the processors of the Accused Products enqueue the network traffic received from either of the first or the second network interfaces on a queue corresponding to the belonging bandwidth zone (e.g., on a queue for one of "Lobby", "VIP Guests", "Meeting Room" in the example on page 13). This page demonstrates that the Accused Products satisfy paragraph g) of claim 1 of the 640 Patent.

81. Page 13 of the User Guide describes how each class's network traffic is processed in a group separate from traffic in other classes in a class-based queuing system in order to achieve different throughput outputs for each class. Page 358 of the User Guide further describes a round robin queuing algorithm that services each queue in turn. Likewise, page 356 of the User Guide explains that packets contain destination addresses. On information and belief, the processors of the Accused Products cycle though the plurality of the queues in order to dequeue the network traffic from the queues and pass the data to one or more outgoing network interfaces to a destination network address. The chart on page 13 of the User Guide shows how network traffic throughput is formed by network traffic from the three different queues (e.g., for each of Lobby, VIP Guests, and Meeting Room). On information and belief, the data is removed data from each class's queue and be passed to the output network interface(s) for transmission to a

destination network address on the Internet and/or guest network. These pages demonstrate that the Accused Products satisfy paragraph h) of claim 1 of the 640 Patent.

82. Page 14 of the User Guide describes how the Accused Products automatically control the effective throughput of the various classes (e.g., class 1 v. class 2). The amount of data dequeued from a particular queue and sent to the network interface controls the effective bandwidth throughput for that queue. Since the throughput is automatically controlled by the Accused Products to meet the desired bandwidth limit for each class, on information and belief, the Accused Products automatically determine how much data to dequeue from that class's queue according to the desired bandwidth limit. This page of the User Guide demonstrates that, when dequeing the network traffic from a particular queue, the one or more processors of the Accused Products are operable to automatically determine an amount of the network traffic to dequeue from the particular queue according to a bandwidth limit for the particular queue, and therefore the Accused Products satisfy paragraph i) of claim 1 of the 640 Patent.

83. Page 104 of the User Guide shows how each of the plurality of first level guaranteed bandwidth zones (PriorityOne, PriorityTwo, PriorityThree) all have a guaranteed bandwidth rate (i.e., they each have a non-zero "Min" bandwidth). Testing the class-based queuing configuration screen on the Test Unit shows that the bandwidth limit for each of the plurality of the queues corresponding to the plurality of the first level guaranteed bandwidth zones (PriorityOne, PriorityTwo, PriorityThree) includes a guaranteed bandwidth rate (e.g., "Min" bandwidths of 25000, 10000, 10000, respectively). This page and the testing reveals how the bandwidth limit for each of the plurality of the queues corresponding to the plurality of the first level guaranteed bandwidth zones in the Accused Products includes a guaranteed bandwidth rate, therefore satisfying paragraph j) of claim 1 of the 640 Patent.

84. Page 103 of the User Guide shows how the fixed total amount of bandwidth of the WAN interface is 50000Kbps while the sum of each of the first level guaranteed bandwidth zones (PriorityOne, PriorityTwo, PriorityThree) is 45000Kbps. The fixed total amount of bandwidth of the one or more first network interfaces (WAN = 50000Kbps) minus the guaranteed bandwidth rate for each of the plurality of the first level guaranteed bandwidth zones (45000Kbps in total) leaves 5000Kbps remaining for the Other bandwidth class. Testing the class-based queuing configuration screen on the Test Unit shows a plurality of first level guaranteed bandwidth zones (PriorityOne, PriorityTwo, PriorityThree) and only one first level remaining bandwidth zone (PriorityFour) not entitled to any guaranteed bandwidth rate. Again, the fixed bandwidth of the WAN 50000Kbps minus each of the first level guaranteed bandwidth zones (45000Kbps in total) leaves 5000Kbps remaining. The first level remaining bandwidth zone (PriorityFour) has no guaranteed bandwidth rate (the "Min" is set to zero) but includes a bandwidth cap equal to 5000Kbps (i.e., the "Max" is set to 5000Kbps). This page and the testing illustrate that the bandwidth limit for the first level remaining bandwidth zone has no guaranteed bandwidth rate but includes a bandwidth cap equal to the fixed total amount of bandwidth of the one or more first network interfaces minus the guaranteed bandwidth rate for each of the plurality of the first level guaranteed bandwidth zones, and therefore the Accused Products satisfy paragraph k) of claim 1 of the 640 Patent.

85. Defendant has directly infringed, and continues to directly infringe, at least claim 1 of the 640 Patent by making, using, selling or offering for sale in the United States, or importing into the United States, the Accused Products, as described above.

86. Defendant has been on actual notice that it infringes the 640 Patent. Guest Tek has properly marked its products. Additionally, Defendant has had actual knowledge of its

infringement at least since February 13, 2018 and February 22, 2018, when Guest Tek advised Defendant by letter of the infringement. Defendant has deliberately, intentionally and willfully infringed the 640 Patent, and continues to do so.

87. As a direct result of Defendant's infringing acts, Guest Tek has suffered and will continue to suffer damage and irreparable harm.

88. Unless Defendant and those acting in active concert with Defendant are enjoined from infringing the 640 Patent, Guest Tek will continue to suffer irreparable injury for which damages are an inadequate remedy.

COUNT IV
INFRINGEMENT OF THE 738 PATENT

89. Guest Tek incorporates each of the preceding paragraphs as if fully set forth herein.

90. The system and method described and claimed in the 738 Patent (Exhibit D) allocates bandwidth among users.

91. Defendant has infringed and is infringing at least claim 20 of the 738 Patent either literally or under the doctrine of equivalents. For ease of reference, claim 20 of the 738 Patent is reproduced below and each paragraph thereof has been annotated with an identifying letter:

> An apparatus for allocating bandwidth between a plurality of bandwidth zones at an establishment serving a plurality of users, each of the plurality of bandwidth zones having a number of users competing for bandwidth allocated thereto, the apparatus comprising:
>
> a) a storage device;
>
> b) a network interface; and
>
> c) a processor coupled to the storage device and the network interface; wherein, by the processor executing software instructions loaded from the storage device, the processor is operable to:

d) provide a first queue and a second queue,

e) wherein the first queue queues first data associated with a first bandwidth zone of the plurality of bandwidth zones, and

f) the second queue queues second data associated with a second bandwidth zone of the plurality of bandwidth zones;

g) repeatedly dequeue a first amount of the first data from the first queue and a second amount of the second data from the second queue, and pass the first amount of the first data and the second amount of the second data to the network interface; and

h) automatically adjust the first amount and the second amount over time

i) such that the first amount is larger than the second amount while the first bandwidth zone has a higher user load than the second bandwidth zone, and

j) such that the first amount is smaller than the second amount while the first bandwidth zone has a lower user load than the second bandwidth zone;

k) wherein user load of the first bandwidth zone is determined according to a number of first users competing for bandwidth in the first bandwidth zone, and

l) user load of the second bandwidth zone is determined according to a number of second users competing for bandwidth in the second zone.

92. Page 11 of the User Guide describes the Accused Products as allocating bandwidth between a plurality of bandwidth zones (e.g., classes) at an establishment (e.g., property) serving a plurality of users. Pages 11 to 12 of the User Guide further describe how each zone (e.g., class) has a number of users competing for bandwidth allocated thereto. These pages of the User Guide reveal how the Accused Products are an apparatus for allocating bandwidth between a plurality of bandwidth zones at an establishment serving a plurality of users, each of

the plurality of bandwidth zones having a number of users competing for bandwidth allocated thereto, and therefore the Accused Products satisfy the preamble of claim 20 of the 738 Patent.

93. On information and belief, the Firmware Upgrade Instructions ("FUI") located at (http://www.nomadix.com/Websites/nomadix/download/agdocs/how-to/admin/NSE-Automated-Firmware-Upgrade-System.pdf, Exhibit G) describes the operation of the Accused Product, in relevant part, insofar as the structure, function and operation of the Patents-In-Suit is concerned.

94. Page 3 of the User Guide discloses that the Accused Products run embedded software called Nomadix Service Engine (NSE). Page 290 of the User Guide discloses that the Accused Products include "system Memory" that can be allocated, free, or in use. Page 1 of the FUI describes how the AG's include a "flash directory" that stores firmware run by the gateway and that can be updated by users. On information and belief, firmware stored in the flash directory is executed by a processor of the Accused products and cause the processor to perform the various actions of the NSE software. The memory and flash directory described by these pages are examples of storage devices provided in the Accused Products. These pages demonstrate that the Accused Products satisfy paragraph a) of claim 20 of the 738 Patent.

95. Page 38 of the User Guide shows that the Accused Products include network interfaces such as WAN, Eth1, Eth2, etc. This page demonstrates how the Accused Products include a network interface and therefore the Accused Products satisfy paragraph b) of claim 20 of the 738 Patent.

96. Page 341 of the User Guide states that an AG is a computing devices that runs software. On information and belief, the Accused Products are computing devices that include a processor coupled to the storage device and the network interface that, by executing software

instructions loaded from the storage device, is operable to perform a plurality of actions, therefore satisfying paragraph c) of claim 20 of the 738 Patent.

97. Page 8 of the User Guide describes that the NSE core functionality includes Class Based Queueing for bandwidth management. Pages 11 to 12 of the User Guide describe a plurality of zones (e.g., classes). Since queuing is performed on a class-basis in class-based queuing and since there are a plurality of classes, it follows that there are a first queue and a second queue provided by the processor of an AG. As shown on pages 11 to 12 of the User Guide, there are three different classes (Conference, Guest Room, Public) that are utilized in a particular example of class-based queuing. Thus, it follows that each of the three classes has its own queue since these three classes are queued on a class basis. These pages of the User Guide show that the Accused Products provide a first queue and a second queue and therefore the Accused Products satisfy paragraph d) of claim 20 of the 738 Patent.

98. Pages 11 to 12 of the User Guide describe a plurality of zones (e.g., classes). There are at least first and second classes (Conference, Guest Room, Public) that are utilized in the example of class-based queuing shown on pages 11 to 12 of the User Guide. On information and belief, a first queue queues first data associated with a first bandwidth zone (e.g., for the Conference zone). These pages of the User Guide show how, in the Accused Products, the first queue queues first data associated with a first bandwidth zone of the plurality of bandwidth zones, and therefore the Accused Products satisfy paragraph e) of claim 20 of the 738 Patent.

99. Pages 11 to 12 of the User Guide describe a plurality of zones (e.g., classes). There are at least first and second classes (Conference, Guest Room, Public) that are utilized in the example of class-based queuing shown on pages 11 to 12 of the User Guide. On information and belief, a second queue queues second data associated with a second bandwidth zone (e.g., for

the Guest room zone). These pages of the User Guide show how, in the Accused Products, the second queue queues second data associated with a second bandwidth zone of the plurality of bandwidth zones, and therefore the Accused Products satisfy paragraph f) of claim 20 of the 738 Patent.

100. Page 13 of the User Guide describes how each class's network traffic is processed in a group separate from traffic in other classes in a class-based queuing system in order to achieve different throughput outputs for each class. Page 358 of the User Guide further describes a round robin queuing algorithm that repeatedly services each queue in turn. The amount of data dequeued from a particular queue and sent to the network interface controls the effective bandwidth throughput for that queue. Since the throughput is automatically controlled by the Accused Products to meet the desired bandwidth limit for each class, on information and belief, the Accused Products repeatedly dequeue a first amount of the first data from the first queue and a second amount of the second data from the second queue, and pass the first amount of the first data and the second amount of the second data to the network interface, therefore satisfying paragraph g) of claim 20 of the 738 Patent.

101. Pages 12-13 of the User Guide shows how the Accused Products dynamically change the values of bandwidth provided for each class's queue over time. Because each class's effective bandwidth is changed over time, on information and belief, the processor of the Accused Products automatically adjusts the first amount and the second amount over time, and therefore the Accused Products satisfy paragraph h) of claim 20 of the 738 Patent.

102. Page 28 of the User Guide describes how Weighted Fair Queuing allocates bandwidth in proportion to individual or group bandwidth limits. The example then shows that when there are 100 users each having 1M individual caps in a first group competing for

bandwidth with 100 users each having 2M individual caps in a second group, the second group gets double the bandwidth than the first group. The ratio is 2:1 because the sum of the 100 * 2M individual caps in the second group is double the 100 * 1M in the first group. The first and second groups equate to first and second bandwidth zones. The ratio of effective bandwidths of different zones corresponds to the ratio of the amounts of data that are being dequeued from the queues associated with those zones. The effective throughput bandwidths outputted are controlled by the Accused Products according to the zone's user load. Furthermore, page 2 of the Nomadix Press Release announcing the release of the "Share unused" bandwidth feature that works in conjunction with class-based queuing mentions how the bandwidth allocations (and therefore the amounts of data that are dequeued) are proportionally increased. The throughput simulator on the Test Unit shows that the effective throughput of a "Plan A" class is higher than the effective throughput of an "Other" class when the load on Plan A is higher than the load on the Other class. Likewise, testing of the Test Unit configured with class-based queuing, weighted fair queuing, and share bandwidth all turned on shows that while there are two users logged into a Plan A, a single user under Plan B loses some of their bandwidth. When Plan A has more users than Plan B, the aggregate bandwidth rate of Plan A is higher than Plan B. The situation reverses when Plan B has more users than Plan A. On information and belief, the Accused Products automatically adjust the first and second amounts to thereby control the throughput speeds of both Plan A and Plan B as the user loads in the two classes change over time. On information and belief, the first amount (achieving 200Kbps) is larger than the second amount (achieving 133Kbps) while the first bandwidth zone (Plan A) has a higher user load (2 users for Plan A v. 1 user for Plan B) than the second bandwidth zone (Plan B). This page and testing reveal that the functionality of the Accused Products adjusts the first and second amounts such that the first

amount is larger than the second amount while the first bandwidth zone has a higher user load than the second bandwidth zone, and therefore satisfy paragraph i) of claim 20 of the 738 Patent.

103. The throughput simulator on the Test Unit shows that the effective throughput of Plan A is lower than the effective throughput of the Other class when the load on Plan A is lower than the load on the Other class. Likewise, testing Test Unit configured with class-based queuing, weighted fair queuing, and share bandwidth all turned on shows that while there are two users logged into a Plan A, a single user under Plan B loses some of their bandwidth. When Plan A has more users than Plan B, the aggregate bandwidth rate of Plan A is higher than Plan B. The situation reverses when Plan B has more users than Plan A. On information and belief, the Accused Products automatically adjust the first and second amounts to thereby control the throughput speeds of both Plan A and Plan B as the user loads in the two classes change over time. On information and belief, the first amount (achieving 133Kbps) is smaller than the second amount (achieving 200Kbps) while the first bandwidth zone (Plan A) has a lower user load (1 user for Plan A v. 2 users for Plan B) than the second bandwidth zone (Plan B). This testing reveals that the Accused Products adjust the first and second amounts such that the first amount is smaller than the second amount while the first bandwidth zone has a lower user load than the second bandwidth zone, therefore satisfying paragraph j) of claim 20 of the 738 Patent.

104. As described above, testing on the Test Unit shows that the user load of the Plan A class depends upon how many user(s) are logged into Plan A and downloading the file under Plan A. Testing on the Test Unit shows the user(s) under Plan A are competing for bandwidth available under Plan A. Page 103 of the User Guide further confirms that the bandwidth limits (Min/Max) are applied across all users in a class, which means that all users in a class are competing for bandwidth in the class. This page and testing confirms that the user load of the

first bandwidth zone is determined by the Accused Products according to a number of first users competing for bandwidth in the first bandwidth zone, therefore satisfying paragraph k) of claim 20 of the 738 Patent.

105. As described above, testing on the Test Unit shows the user load of the Plan B class depends upon how many user(s) are logged into Plan B and downloading the file under Plan B. Testing on the Test Unit shows the user(s) under Plan B are competing for bandwidth available under Plan B. Page 103 of the User Guide further confirms that the bandwidth limits (Min/Max) are applied across all users in a class, which means that all users in a class are competing for bandwidth in the class. This testing confirms that the user load of the second bandwidth zone is determined by the Accused Products according to a number of second users competing for bandwidth in the second bandwidth zone, therefore satisfying paragraph l) of claim 20 of the 738 Patent.

106. Defendant has directly infringed, and continues to directly infringe, at least claim 20 of the 738 Patent by making, using, selling or offering for sale in the United States, or importing into the United States, the Accused Products, as described above.

107. Defendant has been on actual notice that it infringes the 738 Patent. Guest Tek has properly marked its products. Additionally, Defendant has had actual knowledge of its infringement at least since February 13, 2018 and February 22, 2018, when Guest Tek advised Defendant by letter of its infringement of three patents related to the 738 Patent, namely, the asserted 184 Patent, 435 Patent, and 640 Patent, all of which are in the same patent family. Defendant has deliberately, intentionally and willfully infringed the 640 Patent, and continues to do so.

108. As a direct result of Defendant's infringing acts, Guest Tek has suffered and will continue to suffer damage and irreparable harm.

109. Unless Defendant and those acting in active concert with Defendant are enjoined from infringing the 738 Patent, Guest Tek will continue to suffer irreparable injury for which damages are an inadequate remedy.

PRAYER FOR RELIEF

Guest Tek respectfully requests the following relief from this Court:

(a) Judgment that Nomadix has infringed one or more claims of each of the Patents-in-Suit;

(b) Entry of a preliminary and permanent injunction against Nomadix, its officers, agents, servants, employees, attorneys, those in privity with it and those acting in concert with it from further infringement of the Patents-in-Suit, such as through the using, making, selling, or offering for sale in the United States, or importing into the United States, the Accused Products.

(c) An award to Guest Tek of damages for Nomadix's infringement of the Patents-In-Suit, including damages pursuant to 35 U.S.C. § 284, adequate to compensate it for all infringement occurring through the date of judgment, with prejudgment interest, and for any supplemental damages as appropriate, ongoing royalties in the absence of a permanent injunction, and post-judgment interest after that date;

(d) A finding that Nomadix has willfully infringed and is willfully infringing one or more claims of each of the Patents-In-Suit, and awarding treble damages due to Nomadix's deliberate and willful conduct;

(e) A finding that this is an exceptional case and order Nomadix to pay Guest Tek's costs of suit and attorneys' fees pursuant to 35 U.S.C. § 285; and

(f) An award of such other and further relief as the Court may deem just and proper.

DEMAND FOR TRIAL BY JURY

Pursuant to Federal Rule of Civil Procedure 38(b), Guest Tek hereby demands a trial by jury of all issues so triable.

MORRIS, NICHOLS, ARSHT & TUNNELL LLP

/s/ Jack B. Blumenfeld
Jack B. Blumenfeld (#1014)
Jennifer Ying (#5550)
1201 North Market Street
P.O. Box 1347
Wilmington, DE 19899
(302) 658-9200
jblumenfeld@mnat.com
jying@mnat.com

Attorneys for Guest Tek Interactive Entertainment Ltd.

OF COUNSEL:

Steve Rocci
Michael J. Swope
BAKER & HOSTETLER LLP
2929 Arch Street
Cira Centre, 12th Floor
Philadelphia, PA 19104
(215) 568-3100

September 7, 2018
12184574